CLAIMS

We claim:

1. A method for scheduling static and dynamic traffic through a switch fabric including one or more switch slices, comprising for individual such switch slices:

scheduling static traffic by reserving time slots for transmitting said static traffic to at least one destination through a switch slice; and

scheduling dynamic traffic so as not to be transmitting said dynamic traffic to said at least one destination during said reserved time slots through said switch slice.

2. The method according to claim 1, wherein said static traffic includes TDM traffic.

3. The method according to claim 1, wherein said dynamic traffic includes ATM traffic.

4. The method according to claim 1, wherein said dynamic traffic includes IP traffic.

5. The method according to claim 1, wherein said scheduling static traffic includes receiving static traffic information, and storing said static traffic information in a memory of said switch slice.

6. The method according to claim 5, wherein said static traffic information is received at a system configuration time.

7. The method according to claim 5, wherein said static traffic information is received at a run time.

8. The method according to claim 5, wherein said static traffic information includes source and destination information for said static traffic.

9. The method according to claim 8, wherein said static traffic information includes channel information for said static traffic.

10. The method according to claim 1, wherein said scheduling dynamic traffic includes giving a scheduling preference to said static traffic over said dynamic traffic.

11. The method according to claim 10, wherein said scheduling preference is implemented by giving said static traffic a higher relative weight than said dynamic traffic.

12. The method according to claim 1, wherein said scheduling dynamic traffic includes receiving dynamic traffic scheduling requests, aging previously received and not granted dynamic traffic scheduling requests, and discarding expired dynamic traffic scheduling requests to define active dynamic traffic scheduling requests.

13. The method according to claim 12, wherein said scheduling dynamic traffic further includes determining a set of destination winners including a destination winning dynamic traffic scheduling request for each available destination having at least one active dynamic traffic scheduling request for that available destination.

14. The method according to claim 13, wherein said destination winning dynamic traffic scheduling request is determined by selecting a dynamic traffic scheduling request having a highest priority among the at least one active dynamic traffic scheduling request having the same destination.

15. The method according to claim 14, wherein if at least two dynamic traffic scheduling requests for the same destination are tied with the highest priority, then said winning dynamic traffic scheduling request is determined by selecting the dynamic traffic scheduling request having a highest age among said at least two dynamic traffic scheduling requests.

16. The method according to claim 13, wherein said scheduling dynamic traffic further includes determining a set of source winners from said set of destination winners by including a source winning dynamic traffic scheduling request for each source having at least one dynamic traffic scheduling request from that source.

17. The method according to claim 16, wherein said source winning dynamic traffic scheduling request is determined by selecting a dynamic traffic scheduling request having a highest priority among the at least one dynamic traffic scheduling request from the same source.

18. The method according to claim 17, wherein if at least two dynamic traffic scheduling requests from the same source are tied with the highest priority, then said source winning dynamic traffic scheduling request is determined by selecting the dynamic traffic scheduling request having a highest age among said at least two dynamic traffic scheduling requests.

19. The method according to claim 16, wherein if in the process of determining said set of source winners a destination winning dynamic traffic scheduling request is deleted from said set of destination winners, then determining a replacement for said deleted destination winning dynamic traffic scheduling request by determining a new destination winning dynamic traffic scheduling request from among the remaining dynamic traffic scheduling requests having the same destination.

20. An apparatus for scheduling static and dynamic traffic through a switch fabric including one or more switch slices, comprising for individual such switch slices:

a plurality of buffers for storing requests for transmission of dynamic traffic to dynamic traffic destinations through a switch slice;

a memory storing a schedule of static traffic to be transmitted to at least one static traffic destination through said switch slice; and

a grant scheduler coupled to said plurality of buffers and said memory for reserving time slots for transmitting said static traffic to said at least one static traffic destination, and scheduling selected ones of said requests for transmission of dynamic traffic so as not to be transmitting any of said dynamic traffic to said at least one static traffic destination during said reserved time slots through said switch slice.

21. The apparatus according to claim 20, wherein said static traffic includes TDM traffic.

22. The apparatus according to claim 20, wherein said dynamic traffic includes ATM traffic.

23. The apparatus according to claim 20, wherein said dynamic traffic includes IP traffic.

24. The apparatus according to claim 20, wherein said schedule of static traffic is stored in said memory at a system configuration time.

25. The apparatus according to claim 20, wherein said schedule of static traffic is stored in said memory at a run time.

26. The apparatus according to claim 20, wherein said schedule of static traffic includes a plurality of entries individually indicating at least one destination for a given time slot and source.

27. The apparatus according to claim 26, wherein each of said plurality of entries further indicates a channel.

28. The apparatus according to claim 20, wherein said plurality of buffers comprise request shifters that receive dynamic traffic scheduling requests, age previously received and not granted dynamic traffic scheduling requests by shifting them upon each clock pulse controlling transfer of cells to and from said switch slice, and discard expired dynamic traffic scheduling requests by shifting them out of said request shifters.

29. The apparatus according to claim 28, wherein said request shifters are organized into a set for each source including entries, at least one of said entries indicating a primary request and a secondary request from that source for unicast requests.

30. The apparatus according to claim 28, wherein said request shifters are organized into a set for each source including entries, at least one of said entries indicating a request from that source for a multicast request.

31. The apparatus according to claim 28, wherein said request shifters are organized into a set for each source including entries individually having a priority field for storing a priority and a destination field for storing a destination for each request in that entry.

32. The apparatus according to claim 28, wherein a position of an entry in said request shifters indicates an age of said entry.

33. The apparatus according to claim 28, wherein said request shifters are organized into a set for each source including entries individually having a marker field for indicating winning requests that are available to be granted.

34. The apparatus according to claim 20, wherein said grant scheduler comprises a grant scheduler configured to determine a set of destination winners including a destination winning dynamic traffic scheduling request for each available destination having at least one dynamic traffic scheduling requests for that available destination.

35. The apparatus according to claim 34, wherein said grant scheduler is further configured such that said destination winning dynamic traffic scheduling request is determined by selecting a dynamic traffic scheduling request having a highest priority among the at least one dynamic traffic scheduling request for the same destination.

36. The apparatus according to claim 35, wherein said grant scheduler is further configured such that if at least two dynamic traffic scheduling requests for the same destination are tied with the highest priority, then said winning dynamic traffic scheduling request is determined by selecting the dynamic traffic scheduling request having a highest age among said at least two dynamic traffic scheduling requests.

37. The apparatus according to claim 34, wherein said grant scheduler is further configured such that said scheduling dynamic traffic further includes determining a set of source winners from said set of destination winners by including a source winning dynamic traffic scheduling request for each source having at least one dynamic traffic scheduling request from that source.

38. The apparatus according to claim 37, wherein said grant scheduler is further configured such that said source winning dynamic traffic scheduling request is determined by selecting a dynamic traffic scheduling request having a highest priority among the at least one dynamic traffic scheduling request from the same source.

39. The apparatus according to claim 38, wherein said grant scheduler is further configured such that if at least two dynamic traffic scheduling requests from the same source are tied with the highest priority, then said source winning dynamic traffic scheduling request is determined by selecting the dynamic traffic scheduling request having a highest age among said at least two dynamic traffic scheduling requests.

40. The apparatus according to claim 39, wherein said grant scheduler is further configured such that if in the process of determining said set of source winners a destination winning dynamic traffic scheduling request is deleted from said set of destination winners, then determining a replacement for said deleted destination winning dynamic traffic scheduling request by determining a new destination winning dynamic traffic scheduling request from among the remaining dynamic traffic scheduling requests having the same destination.

41. A method for scheduling dynamic traffic through a switch fabric including one or more switch slices, comprising for individual such switch slices:

(a) receiving a plurality of dynamic traffic scheduling requests individually having an associated priority, source, and destination;

(b) incrementing ages of previously received and ungranted dynamic traffic scheduling requests individually having an associated priority, source, and destination;

(c) generating relative weights for said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests based upon their associated priorities and ages; and

(d) determining a set of dynamic traffic scheduling requests to be granted from said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests using said relative weights such that no two dynamic traffic

scheduling requests in said set has the same associated source or destination.

42. The method according to claim 41, wherein (c) comprises concatenating said associated priority and age for each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests such that said associated priority is positioned so as to be given more weight than said associated age.

43. The method according to claim 42, wherein each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests is either a unicast or multicast request, and (c) comprises concatenating said associated priority, a unicast/multicast indication, and said associated age for each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests such that said associated priority is positioned so as to be given more weight than said unicast/multicast indication and said associated age, and said unicast/multicast indication is positioned so as to be given more weight than said associated age.

44. The method according to claim 43, wherein each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests are either primary or secondary requests, and (c) comprises concatenating said associated priority, said unicast/multicast indication, said associated

age, and a primary/secondary request indication for each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests such that said associated priority, said unicast/multicast indication and said associated age are positioned so as to each be given more weight than said primary/secondary request indication.

45. The method according to claim 43, wherein if two or more dynamic traffic scheduling requests are tied with the highest relative weight compared to other dynamic traffic scheduling requests having the same associated destination, then selecting one of said two or more dynamic traffic scheduling requests to be included in said set of dynamic traffic scheduling requests to be granted, by the following tie-breaking conditions:

if said two or more dynamic traffic scheduling requests are unicast requests, then selecting one of said two or more dynamic traffic scheduling requests that has an associated source that is closest in an ordered sequence with wrap-around after a source from which a request had most recently been granted to the same associated destination; and

if said two or more dynamic traffic scheduling requests are multicast requests, then selecting one of said two or more dynamic traffic scheduling requests that has an associated source that is closest after a source indicated by a global pointer that is advanced by one source in an ordered sequence with wrap-around upon each cell transfer clock cycle of said switch slice.

46. The method according to claim 45, wherein unicast requests are given preference over multicast requests if said switch slice has been receiving more unicast requests than multicast requests, and multicast requests are given preference over unicast requests if said switch slice has been receiving more multicast requests than unicast requests.

47. The method according to claim 46, wherein the preference of unicast or multicasts requests is toggled following each cell transfer clock cycle of said switch slice if said switch slice has been receiving approximately equal numbers of said multicast and unicast requests.

48. The method according to claim 45, wherein unicast requests are given preference over multicast requests if more unicast requests have been expiring due to age than multicast requests, and multicast requests are given preference over unicast requests if more multicast requests have been expiring due to age than unicast requests.

49. The method according to claim 41, wherein (d) comprises:

(d1) performing a destination round-robin tournament wherein each of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests is compared against each other such that if two or more of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling

requests have the same associated destination, then the one that has the highest relative weight shall be declared the winner, and if one or more of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests are the only ones requesting their respective associated destinations, then said one or more of said plurality of dynamic traffic scheduling requests and said previously received and ungranted dynamic traffic scheduling requests shall each be declared a winner; and

(d2) performing a source round-robin tournament wherein each winner of said destination round-robin tournament is compared against each other winner such that if two or more of said winners of said destination round-robin tournament have the same associated source, then the one that has the highest relative weight shall be declared the winner and the others shall not be considered a valid request for the remainder of the cell transfer clock cycle of said switch slice.

50. The method according to claim 49, wherein (d) further comprises:

(d3) if in the process of performing said source round-robin tournament at least one of said winners of said destination round-robin tournament shall not be considered a valid request for the remainder of the cell transfer clock cycle of said switch slice, then determining, if possible, a replacement for said at least one of said winners by repeating (d1) to (d3) until no more winners shall be eliminated as being not considered a valid request for the remainder of the cell transfer clock cycle of said switch slice.

51. An apparatus for scheduling dynamic traffic through a switch fabric including one or more switch slices, comprising for individual such switch slices:

a plurality of buffers storing requests for transmission of dynamic traffic through a switch slice, each of said requests having an associated priority, source, destination, and age; and

a grant scheduler coupled to said plurality of buffers for determining a set of requests to be granted by generating relative weights for said requests based at least in part upon their associated priorities and ages, and determining a set of requests to be granted using said relative weights such that no two requests in said set has the same associated source or destination.

52. The apparatus according to claim 51, wherein said relative weight is generated by concatenating said associated priority and age for each of said requests such that said associated priority is positioned so as to be given more weight than said associated age.

53. The apparatus according to claim 52, wherein individual ones of said plurality of requests are either unicast or multicast requests, and said relative weight is generated by concatenating said associated priority, a unicast/multicast indication, and said associated age for each of said individual ones of said plurality of requests such that said associated priority is positioned so as to be given more weight than said unicast/multicast indication and said associated age, and said unicast/multicast indication

is positioned so as to be given more weight than said associated age.

54. The apparatus according to claim 53, wherein individual ones of said plurality of requests are either a primary or secondary request, and said relative weight is generating by concatenating said associated priority, said unicast/multicast indication, said associated age, and a primary/secondary request indication for each of said plurality of requests such that said associated priority, said unicast/multicast and said associated age are positioned so as to each be given more weight than said primary/secondary request indication.

55. The apparatus according to claim 53, further comprising:

a plurality of pointers associated with corresponding ones of said associated destinations of said plurality of requests, and indicating sources from which requests had last been granted to said corresponding ones of said associated destinations;

wherein if two or more requests are tied with the highest relative weight compared to other requests having the same associated destination and if said two or more requests are unicast requests, then said grant scheduler determines which one of said two or more requests to be included in said set of requests to be granted by selecting the one that has an associated source that is closest in an ordered sequence with wrap-around after the source indicated by the one of said plurality of pointers associated with said same associated destination.

56. The apparatus according to claim 55, further comprising:

a global pointer indicating one of said associated sources resulting from advancing one source at a time in a predetermined repeating sequence of said associated sources upon each cell transfer clock cycle of said switch slice;

wherein if two or more requests are tied with the highest relative weight compared to other requests having the same associated destination and if said two or more requests are multicast requests, then said grant scheduler determines which one of said two or more requests to be included in said set of requests to be granted by selecting the one that has an associated source that is closest in an ordered sequence with wrap-around after the source indicated by said global pointer according to said predetermined repeating sequence of said associated sources.

57. The apparatus according to claim 56, further comprising a configuration register programmable to indicate a preference between unicast and multicast requests, wherein a first value indicates that unicast requests are given preference over multicast requests, a second value indicates that multicast requests are given preference over unicast requests, and a third value indicates that unicast and multicast requests are given equal preference.

58. The apparatus according to claim 57, wherein if said switch slice has been receiving more unicast requests than multicast requests, then said configuration register is programmed to indicate said first value; if said

switch slice has been receiving more multicast requests than unicast requests, then said configuration register is programmed to indicate said second value; and if said switch slice has been receiving approximately equal numbers of unicast and multicast requests, then said configuration register is programmed to indicate said third value.

59. The apparatus according to claim 57, wherein if more unicast requests have been expiring due to age than multicast requests, then said configuration register is programmed to indicate said first value; if more multicast requests have been expiring due to age than unicast requests, then said configuration register is programmed to indicate said second value; and if approximately equal numbers of unicast and multicast requests have been expiring due to age, then said configuration register is programmed to indicate said third value.

60. The apparatus according to claim 57, wherein the preference of unicast or multicasts requests is toggled following each cell transfer clock cycle of said switch slice if said configuration register indicates said third value.

61. The apparatus according to claim 56, further comprising a plurality of configuration registers respectively associated with different time slots and individually programmable to indicate a preference between unicast and multicast requests, wherein a first value indicates that unicast requests are given preference over

multicast requests, a second value indicates that multicast requests are given preference over unicast requests, and a third value indicates that unicast and multicast requests are given equal preference.

62. The apparatus according to claim 51, wherein said grant scheduler determines said set of requests to be granted by:

(d1) performing a destination round-robin tournament wherein each of said plurality of requests is compared against each other such that if two or more of said plurality of requests have the same associated destination, then the one that has the highest relative weight shall be declared the winner, and if one or more of said plurality of requests are the only ones requesting their respective associated destinations, then said one or more of said plurality of requests shall each be declared a winner; and

(d2) performing a source round-robin tournament wherein each winner of said destination round-robin tournament is compared against each other winner such that if two or more of said winners of said destination round-robin tournament have the same associated source, then the one that has the higher relative weight shall be declared the winner and the others shall not be considered a valid request for the remainder of the cell transfer clock cycle of said switch slice.

63. The apparatus according to claim 62, wherein said grant scheduler further determines said set of requests to be granted by:

(d3) if in the process of performing said source round-robin tournament at least one of said winners of said destination round-robin tournament shall not be considered a valid request for the remainder of the cell transfer clock cycle of said switch slice, then determining, if possible, a replacement for said at least one of said winners by repeating (d1) to (d3) until no more winners shall be eliminated as being not considered a valid request for the remainder of the cell transfer clock cycle of said switch slice.

64. A SONET/SDH network element, comprising:

a plurality of line cards individually having a plurality of switch interface transmission ports and a plurality of switch interface receiving ports; and

a plurality of switch slices individually coupled to each of said plurality of line cards, and individually including means for scheduling static traffic from one of said plurality of line cards to another or the same one of said plurality of line cards by reserving time slots for transmitting said static traffic through said individual switch slice, and means for scheduling dynamic traffic so as not to be transmitting said dynamic traffic to said another or the same one of said plurality of line cards during said reserved time slots through said individual switch slice.